

What is claimed is:

- 1 1. A method of determining data placement for a distributed storage system
2 comprising the steps of:
3 selecting a heuristic class which meets a performance requirement and
4 which provides a replication cost that is within an allowable limit of a
5 minimum replication cost; and
6 instantiating a data placement heuristic selected from a range of data
7 placement heuristics according to the heuristic class.
- 1 2. The method of claim 1 wherein the performance requirement comprises a bi-
2 modal performance metric.
- 1 3. The method of claim 2 wherein the bi-modal performance metric comprises a
2 criterion and a ratio of successful requests to total requests.
- 1 4. The method of claim 1 wherein the data placement heuristic comprises a
2 computer implemented technique of placing data objects onto nodes of the
3 distributed storage system.
- 1 5. The method of claim 4 further comprising the step of evaluating a placement
2 of the data objects.
- 1 6. The method of claim 5 wherein the step of evaluating the data placement
2 heuristic provides a performance result and a cost result for the system
3 configuration and the workload.
- 1 7. The method of claim 5 wherein the step of instantiating the data placement
2 heuristic comprises simulating an instantiation of the data placement heuristic.
- 1 8. The method of claim 7 further comprising the steps of:
2 selecting a second heuristic class for the workload and a second system
3 configuration;
4 instantiating a second data placement heuristic according to the second

5 heuristic class; and

6 evaluating a second placement of the data objects made according to

7 the second data placement heuristic.

1 9. The method of claim 7 further comprising the steps of:

2 selecting a second heuristic class for the system configuration and a

3 second workload;

4 instantiating a second data placement heuristic according to the second

5 heuristic class; and

6 evaluating a second placement of the data objects made according to

7 the second data placement heuristic.

1 10. The method of claim 5 wherein the step of instantiating the data placement

2 heuristic comprises instantiating the data placement heuristic on an actual

3 distributed storage system operating with an actual workload.

1 11. The method of claim 10 further comprising the steps of:

2 selecting a second heuristic class for the system configuration and the

3 actual workload;

4 instantiating a second data placement heuristic according to the second

5 heuristic class; and

6 evaluating a second placement of the data objects made according to

7 the second data placement heuristic.

1 12. The method of claim 1 wherein the performance requirement comprises a data

2 access latency.

1 13. The method of claim 1 wherein the performance requirement comprises an

2 average data access latency.

1 14. The method of claim 1 wherein the performance requirement comprises a data

2 access bandwidth.

1 15. The method of claim 1 wherein the performance requirement comprises a data

2 update time.

1 16. The method of claim 1 wherein the step of selecting the heuristic class
2 determines a plurality of heuristic parameters.

1 17. The method of claim 16 wherein the step of instantiating the data placement
2 heuristic instantiates the data placement heuristic according to the heuristic
3 parameters.

1 18. The method of claim 17 wherein the step of instantiating the data placement
2 heuristic sets other heuristic parameters to defaults.

1 19. The method of claim 1 wherein the replication cost comprises data storage
2 cost.

1 20. The method of claim 1 wherein the replication cost comprises a replica
2 creation cost.

1 21. The method of claim 20 wherein the replication creation cost comprises a
2 network bandwidth cost for transferring replicas and replica changes.

1 22. The method of claim 20 wherein the replica creation cost comprises a system
2 load cost for running the data placement heuristic.

1 23. A method of determining data placement for a distributed storage system
2 comprising the steps of:
3 selecting a heuristic class which meets a performance requirement and
4 which provides a replication cost that is within an allowable limit of a
5 minimum replication cost;
6 instantiating a data placement heuristic selected from a range of data
7 placement heuristics according to the heuristic class; and
8 evaluating a placement of data objects onto nodes of the distributed
9 storage system made according to the data placement heuristic.

1 24. The method of claim 23 wherein the step of instantiating the data placement
2 heuristic comprises simulating instantiation of the data placement heuristic.

1 25. The method of claim 23 wherein the step of instantiating the data placement
2 heuristic comprises instantiating the data placement heuristic on an actual
3 distributed storage system operating with an actual workload.

1 26. A method of determining data placement for a distributed storage system
2 comprising the steps of:
3 selecting a heuristic class which meets a performance requirement and
4 which provides a replication cost that is within an allowable limit of a
5 minimum replication cost;
6 instantiating a data placement heuristic selected from a range of data
7 placement heuristics according to the heuristic class;
8 evaluating a placement of data objects onto nodes of the distributed
9 storage system made according to the data placement heuristic; and
10 iteratively performing the steps of selecting the heuristic class,
11 instantiating the data placement heuristic, and evaluating the placement of
12 the data objects.

1 27. The method of claim 26 wherein second and subsequent performance of the
2 steps of selecting the heuristic class, instantiating the data placement heuristic, and
3 evaluating the placement of the data objects seeks to improve the data placement
4 heuristic.

1 28. The method of claim 26 wherein second and subsequent performance of the
2 steps of selecting the heuristic class, instantiating the data placement heuristic, and
3 evaluating the placement of the data objects seeks to modify the data placement
4 heuristic to account for a changing workload.

1 29. A computer readable memory comprising computer code for implementing a
2 method of determining data placement for a distributed storage system, the
3 method of determining the data placement comprising the steps of:
4 selecting a heuristic class which meets a performance requirement and

5 which provides a replication cost that is within an allowable limit of a
6 minimum replication cost; and
7 instantiating a data placement heuristic selected from a range of data
8 placement heuristics according to the heuristic class.

1 30. A computer readable memory comprising computer code for implementing a
2 method of determining data placement for a distributed storage system, the
3 method of determining the data placement comprising the steps of:
4 selecting a heuristic class which meets a performance requirement and
5 which provides a replication cost that is within an allowable limit of a
6 minimum replication cost;
7 instantiating a data placement heuristic selected from a range of data
8 placement heuristics according to the heuristic class; and
9 evaluating a placement of data objects onto nodes of the distributed
10 storage system made according to the data placement heuristic.

1 31. A computer readable memory comprising computer code for implementing a
2 method of determining data placement for a distributed storage system, the
3 method of determining the data placement comprising the steps of:
4 selecting a heuristic class which meets a performance requirement and
5 which provides a replication cost that is within an allowable limit of a
6 minimum replication cost;
7 instantiating a data placement heuristic selected from a range of data
8 placement heuristics according to the heuristic class;
9 evaluating a placement of data objects onto nodes of the distributed
10 storage system made according to the data placement heuristic; and
11 iteratively performing the steps of selecting the heuristic class,
12 instantiating the data placement heuristic, and evaluating the placement of
13 the data objects.